arXiv:1610.08568v1 [math.OC] 26 Oct 2016

Stochastic First-Order Minimization Techniques Using Jensen Surrogates for X-Ray Transmission Tomography

Soysal Degirmenci[a], Joseph A. O'Sullivan[a], David G. Politte[b]

[a]Department of Electrical and Systems Engineering, Washington University,
1 Brookings Drive, Saint Louis, MO, USA, 63130
[b]Mallinckrodt Institute of Radiology, Washington University School of Medicine,
510 South Kingshighway Blvd, Saint Louis, MO, USA, 63110

ABSTRACT

Image reconstruction in X-ray transmission tomography has been an important research field for decades. In light of data volume increasing faster than processor speeds, one needs accelerated iterative algorithms to solve the optimization problem in the X-ray CT application. Incremental methods, in which a subset of data is being used at each iteration to accelerate the computations, have been getting more popular lately in the machine learning and mathematical optimization fields. The most popular member of this family of algorithms in the X-ray CT field is ordered-subsets. Even though it performs well in earlier iterations, the lack of convergence in later iterations is a known phenomenon. In this paper, we propose two incremental methods that use Jensen surrogates for the X-ray CT application, one stochastic and one ordered-subsets type. Using measured data, we show that the stochastic variant we propose outperforms other algorithms, including the gradient descent counterparts.

Keywords: X-ray imaging, stochastic gradient descent, ordered subsets, Jensen surrogates, incremental descent, stochastic average

1. INTRODUCTION

With the scale of computational inverse problems getting larger in many applications, including X-Ray Computed Tomography, acceleration techniques that provide fast rates of convergence have become crucial. One of the acceleration techniques widely used in many applications is incremental methods, which is a range decomposition technique.[1] In the general case, one forms a surrogate function around the current estimate using a subset of data and computes the next estimate by minimizing it. The idea behind the acceleration is that the subset of data is a good approximation of the full data space and computational time is saved by using only a portion of it. The most popular variants of incremental methods, ordered subsets, and stochastic surrogate descent, perform well in earlier iterations but are known to lack convergence for the later iterations. Lately, the variants of incremental methods in optimization field have proliferated. The algorithms proposed were shown to perform well with desirable convergence properties. The idea behind these methods is to keep track of the surrogate parameters for each subset and to use a combination of them while performing updates, rather than using only one of them.

Jensen surrogates is a possible choice of surrogate functions, and it has nice properties that result in algorithms with attractive properties and good performance.[2,3] For X-ray CT, the family of algorithms that uses Jensen surrogates has parameters that are easy to compute. These can also be extended to acceleration methods, including range-based techniques.

In this paper, we:

Soysal Degirmenci: E-mail: s.degirmenci@wustl.edu
Joseph A. O'Sullivan: E-mail: jao@wustl.edu
David G. Politte: E-mail: politted@wustl.edu

- explore stochastic incremental methods and propose two novel accelerated algorithms using Jensen surrogates for X-ray CT, Stochastic Average Jensen Surrogates Optimization (SA-JS) and Ordered Subsets Average Jensen Surrogates Optimization (OSA-JS), and
- show that the SA-JS algorithm outperforms the other competing algorithms for a large number of subsets by using two sets of real data collected from a baggage scanner.

2. CONVEX OPTIMIZATION USING JENSEN SURROGATES FOR X-RAY CT

The objective function we minimize for monoenergetic X-Ray transmission tomography is as follows.

$$\min_{\boldsymbol{x}\in\mathbb{X}} \Phi(\boldsymbol{x}) = \min_{\boldsymbol{x}\in\mathbb{X}} f(\boldsymbol{x}) + \lambda\beta(\boldsymbol{x}) \quad = \quad \min_{\boldsymbol{x}\in\mathbb{X}} \sum_{i=1}^{M} f_i(\boldsymbol{x}) + \lambda \sum_{k=1}^{K} \beta_k(\boldsymbol{x}) \tag{1}$$

$$= \min_{\boldsymbol{x}\in\mathbb{R}_+^N} \sum_{i=1}^{M} \tilde{f}_i((\boldsymbol{H}\boldsymbol{x})_i) + \lambda \sum_{k=1}^{K} \tilde{\beta}_k((\boldsymbol{C}\boldsymbol{x})_k), \tag{2}$$

where $f(\boldsymbol{x})$ is the data-fitting term and $\beta(\boldsymbol{x})$ is the regularization term. Here, we specifically look at the Poisson log-likelihood data-fitting term, where

$$\tilde{f}_i(l) = d_i l + I_{0,i}\exp(-l), \tag{3}$$

$\boldsymbol{d} \in \mathbb{R}_+^M$ is the attenuated data vector, $\boldsymbol{I_0} \in \mathbb{R}_+^M$ is the incident photon count vector, and $\boldsymbol{H} \in \mathbb{R}_+^{M\times N}$ is the system matrix that defines the relationship between ray-paths and voxels. In a simple ray-tracing model, h_{ij} represents the length of intersection between the voxel indexed by j and the ray-path indexed by i. The non-negativity constraint on the image is due to the physical nature of linear attenuation coefficients. The regularization term we use is

$$\tilde{\beta}_k(t) = \omega_k \delta^2 \left(\left|\frac{t}{\delta}\right| - \log(1 + \left|\frac{t}{\delta}\right|) \right), \tag{4}$$

where $\omega_k > 0$, $\delta > 0$, and $\boldsymbol{C} \in \mathbb{R}^{K\times N}$ is a matrix that has 1s along the diagonal and only one -1 off-diagonal for each row. We further assume that for two arbitrary voxels, if there exists a row in $\boldsymbol{C}$ where the first voxel has a value equal to 1 and the second -1, there should exist another row where the reverse is true as well. The $\tilde{\beta}_k$ defined is an edge-preserving function that is of Huber type, convex, even, and differentiable, where $\boldsymbol{C}$ defines a neighborhood around a center voxel we would like to use to regularize the image. Similarly, ω_k determines the weights of the corresponding neighborhood. This function behaves quadratically for small $|t/\delta|$ and linearly for large $|t/\delta|$ values. Since both $\tilde{f}_i$ and $\tilde{\beta}_k$ are convex and differentiable, we can use the same Jensen surrogate formulation for each of them to create an iterative minimization algorithm.

For the data-fitting term, $f(\boldsymbol{x})$, we denote its Jensen surrogate around $\boldsymbol{x}$ parameterized by $\boldsymbol{r}$ as $g_{\boldsymbol{r}}(\boldsymbol{x};\hat{\boldsymbol{x}})$. We follow the same procedure as in,[3] which gives the resulting surrogate*

$$g_{\boldsymbol{r}}(\boldsymbol{x};\hat{\boldsymbol{x}}) = \sum_i g_{i,\boldsymbol{r}_i}(\boldsymbol{x};\hat{\boldsymbol{x}}) \quad = \quad \sum_i \sum_j r_{ij} \tilde{f}_i\Big(\frac{h_{ij}}{r_{ij}}(x_j - \hat{x}_j) + (\boldsymbol{H}\hat{\boldsymbol{x}})_i\Big) + \text{const.}$$

$$= \sum_i \sum_j h_{ij} d_i (x_j - \hat{x}_j) + r_{ij} I_{0,i} \exp(-\frac{h_{ij}}{r_{ij}}(x_j - \hat{x}_j) - (\boldsymbol{H}\hat{\boldsymbol{x}})_i). \tag{5}$$

*We ignore the constant term after step 1.

We choose $r_{ij} = h_{ij}/Z$, where $Z = \max_i \sum_j h_{ij}$, which results in a surrogate function as follows.

$$\begin{aligned} g_{\boldsymbol{r}}(\boldsymbol{x};\hat{\boldsymbol{x}}) = \sum_i g_{i,\boldsymbol{r}_i}(\boldsymbol{x};\hat{\boldsymbol{x}}) &= \sum_i \sum_j \frac{h_{ij}}{Z} \tilde{f}_i\Big(Z(x_j - \hat{x}_j) + (\boldsymbol{H}\hat{\boldsymbol{x}})_i\Big) + \text{const.} \\ &= \sum_i \sum_j h_{ij} d_i(x_j - \hat{x}_j) + \frac{h_{ij}}{Z} I_{0,i} \exp(-Z(x_j - \hat{x}_j) - (\boldsymbol{H}\hat{\boldsymbol{x}})_i) \\ &= \sum_i \sum_j h_{ij} d_i(x_j - \hat{x}_j) + \frac{h_{ij}}{Z} \hat{q}_i \exp(-Z(x_j - \hat{x}_j)) \\ &= \sum_i b_j(x_j - \hat{x}_j) + \frac{\hat{b}_j}{Z} \exp(-Z(x_j - \hat{x}_j)). \end{aligned} \tag{6}$$

Here, it is important to note that when $r_{ij} = \frac{h_{ij}\hat{x}_j}{\sum_{j'} h_{ij'}\hat{x}_{j'}}$, the surrogate does not have a closed-form minimization and thus must be solved using some convex minimization method. Lange, Fessler et al.[4] use this auxiliary variable choice and solves it via Newton's method.

For the regularization term $\beta(\boldsymbol{x})$ we denote its Jensen surrogate around $\boldsymbol{x}$ parameterized by $\boldsymbol{s}$ as $B_{\boldsymbol{s}}(\boldsymbol{x};\hat{\boldsymbol{x}})$. We choose $s_{ij} = |c_{ij}|/2$ (since $c_{ij} \in \{-1, 0, 1\}$ and there is exactly one element equal to 1 and one element equal to -1 in each row of $\boldsymbol{C}$, the denominator is equal to 2). Then the surrogate becomes[†]

$$\begin{aligned} B_{\boldsymbol{s}}(\boldsymbol{x};\hat{\boldsymbol{x}}) = \sum_k B_{k,\boldsymbol{s}}(\boldsymbol{x};\hat{\boldsymbol{x}}) &= \sum_k \sum_j s_{kj} \tilde{\beta}_k\Big(\frac{c_{kj}}{s_{kj}}(x_j - \hat{x}_j) + (\boldsymbol{C}\hat{\boldsymbol{x}})_k\Big) + \text{const.} \\ &= \sum_k \sum_j \frac{|c_{kj}|}{2} \tilde{\beta}_k\Big(2\,\text{sgn}\,(c_{kj})(x_j - \hat{x}_j) + (\boldsymbol{C}\hat{\boldsymbol{x}})_k\Big) \\ &= \sum_{j' \in N_j} \sum_j \frac{|c_{j'j}|}{2} \tilde{\beta}_{j'}\Big(2\,\text{sgn}\,(c_{j'j})(x_j - \hat{x}_j) + \text{sgn}\,(c_{j'j})(\hat{x}_j - \hat{x}_{j'})\Big) \\ &= \sum_j \sum_{j' \in N_j} \frac{|c_{j'j}|}{2} \tilde{\beta}_{jj'}\Big(2\,\text{sgn}\,(c_{j'j})(x_j - \hat{x}_j) + \text{sgn}\,(c_{j'j})(\hat{x}_j - \hat{x}_{j'})\Big) \\ &= \sum_j \sum_{j' \in N_j} \frac{1}{2} \tilde{\beta}_{jj'}\Big(2(x_j - \hat{x}_j) + (\hat{x}_j - \hat{x}_{j'})\Big) \\ &= \sum_j \sum_{j' \in N_j} \frac{1}{2} \tilde{\beta}_{jj'}\Big(2x_j - \hat{x}_j - \hat{x}_{j'}\Big) \\ &= \sum_j \sum_{j' \in N_j} \frac{\omega_{jj'}}{2} \delta^2 \Bigg(\Big|\frac{2x_j - \hat{x}_j - \hat{x}_{j'}}{\delta}\Big| - \\ &\qquad \log(1 + \Big|\frac{2x_j - \hat{x}_j - \hat{x}_{j'}}{\delta}\Big|)\Bigg), \end{aligned} \tag{7}$$

where in the third step we changed the notation, so that N_j is a set of indices that defines the neighborhood of index j (those sets of indices are the only ones that are relevant to x_j.), and in the fifth step we used the fact that $\tilde{\beta}$ is an even function.

When two surrogate functions are combined, we have our algorithm, the Jensen Surrogates Optimization for X-Ray CT, presented in Algorithm 1.

[†]We ignore the constant term after step 1.

Algorithm 1 Jensen Surrogates Optimization for X-Ray CT (Full-JS)

Input : $\boldsymbol{x}^{(0)} \in \mathbb{R}_+^N, \boldsymbol{d}, \boldsymbol{I_0} \in \mathbb{R}_+^M, \boldsymbol{H} \in \mathbb{R}^{M \times N}, \lambda \geq 0, \delta > 0.$
Pre-compute $b_j = \sum_i d_i h_{ij}, \forall j.$
Pre-compute $Z = \max_i \sum_j h_{ij}.$
for $n = 0, 1, 2, ...$ **do**
  $q_i^{(n)} = I_{0,i} \exp(-\sum_j h_{ij} x_j^{(n)}), \forall i.$
  $b_j^{(n)} = \sum_i q_i^{(n)} h_{ij}, \forall j.$
  $x_j^{(n+1)} = \operatorname{argmin}_{x \geq 0} b_j(x - x_j^{(n)}) + b_j^{(n)}/Z \exp(-Z(x - x_j^{(n)})) + \lambda \sum_{j' \in N_j} \frac{1}{2} \tilde{\beta}_{jj'} (2x - \hat{x}_j - \hat{x}_{j'}), \forall j.$
end

As in most iterative algorithms, this algorithm consists of three steps in each iteration. First, the current estimate in the image domain is forward projected to the data domain with Beer's Law of attenuation applied. Then, this estimate is back projected onto the image domain where this information as well as the back projection of attenuated data and the previous image estimate are used to compute the next estimate. In the case of no regularization ($\lambda = 0$), the surrogate can be minimized in one step; there is a closed-form update. For the regularized case, however, there is not a closed-form update that minimizes the combined surrogate functions. Instead, we have N independent one-dimensional convex problems we can minimize in parallel. Any convex optimization method can be used to minimize these functions. In order to achieve fast convergence we first attempted to use Newton's method. It is important to note that since each minimization is a one-dimensional problem, inversion of the Hessian is not an issue and we take advantage of that fact. However, due to characteristics of the functions being minimized, Newton's method diverged for some cases. Then, we attempted to use a Trust Region Newton's method,[5] which is a modification of Newton's method. In the trust region method, in each iteration, there is a trust region defined such that it bounds the next iterate computed, which makes the algorithm more stable. A metric that measures how well the quadratic approximation approximates the original function is computed at each iteration. Depending on the value of this metric, we either "trust" the quadratic approximation more and expand the trust region, or trust less and shrink it. Unfortunately, the trust region method requires two computations of the function, one first derivative, one second derivative, and many comparisons per iteration. In order to get a faster method that requires fewer computations per iteration and produces comparable performance, we developed a modified trust region method. This method takes advantage of the structure of the problem to construct a fixed trust region and only requires one first derivative, one second derivative, and two comparisons per iteration. More information about these methods can be found in Degirmenci.[6]

As discussed above, when $\lambda = 0$, the image update at iteration n becomes

$$x_j^{(n+1)} = \left[x_j^{(n)} - \frac{1}{Z} \log\left(\frac{b_j}{b_j^{(n)}} \right) \right]_+, \forall j, \tag{8}$$

where $[\cdot]_+ = \max(0, \cdot)$ is a non-negativity operator. Unfortunately, these types of "full" methods have a sublinear rate of convergence,[6] which is very slow. Applications of X-ray imaging, such as baggage scanning (considered below) and medical imaging, both require fast reconstructions with "good" resultant image volumes. In the next section, we will look at several acceleration methods based on range decompositions, which are variants of Algorithm 1.

3. ACCELERATION METHODS

When the amount of data is very large, a reasonable approach to reach the optimal solution faster is to use a portion or subset of the data at each iteration to perform updates. These methods are called "incremental methods" in general. In this section, we first present the popular incremental approach in X-ray CT called ordered subsets using Jensen surrogates and propose our new algorithms Stochastic Average and Ordered Subsets Average Jensen Surrogates Optimization for X-Ray CT. For convenience, we assume that the data is split into B^r subsets where the set of source-detector indices of subset k is represented as $\mathcal{B}_k^r$.

3.1 Ordered Subsets Jensen Surrogates Optimization for X-Ray CT

Ordered subsets is a popular technique used in X-ray CT in which the estimates are updated at each iteration by using a subset of data in an ordered way. In other words, the choice of indices to be used is deterministic and follows a cyclic order. Since only a subset of data is used, in order to keep the balance between the data-fitting and the regularization term, the regularization parameter is scaled down by B^r, the number of subsets. The corresponding algorithm of ordered subsets when used with Jensen surrogates is presented in Algorithm 2. This

Algorithm 2 Ordered Subsets Jensen Surrogates Optimization for X-Ray CT (OS-JS)

Input : $\boldsymbol{x}^{(0)} \in \mathbb{R}_+^N, \boldsymbol{d}, \boldsymbol{I_0} \in \mathbb{R}_+^M, \boldsymbol{H} \in \mathbb{R}^{M \times N}, \lambda \geq 0, \delta > 0, \mathcal{B}_k^r$ for $k = 0, 1, ..., (B^r - 1)$
Pre-compute $b_j^k = \sum_{i \in \mathcal{B}_k^r} d_i h_{ij}, \forall j, k.$
Pre-compute $Z = \max_i \sum_j h_{ij}.$
for $n = 0, 1, 2, ...$ **do**
  $k = \mod(n, B^r)$
  $q_i^{(n)} = I_{0,i} \exp(-\sum_j h_{ij} x_j^{(n)}), \forall i \in \mathcal{B}_k^r.$
  $b_j^{(n,k)} = \sum_{i \in \mathcal{B}_k^r} q_i^{(n)} h_{ij}, \forall j.$
  $x_j^{(n+1)} = \operatorname{argmin}_{x \geq 0} b_j^k (x - x_j^{(n)}) + b_j^{(n,k)} / Z \exp(-Z(x - x_j^{(n)})) + \frac{\lambda}{B^r} B_{\boldsymbol{s}}(\boldsymbol{x}; \boldsymbol{x}^{(n)}), \forall j.$
end

algorithm is known to work well in earlier iterations when the estimate is "far away" from the optimum. However, for later iterations, it lacks convergence and stops decreasing after some number of iterations. For the gradient descent case, Bertsekas et al.[1] showed that regardless of the number of iterations, this type of method never gets closer to the optimal function value than a certain positive constant.

3.2 Stochastic Average Jensen Surrogates Optimization for X-Ray CT

As seen in the previous section, when the ordered subsets technique is used, we use only the most recent version of the back projection information in order to minimize a subset of the problem. In this section, we propose to store the back projection of the estimates for each subset, whether they are the most recent or not, and to use a combination of this back projection information to update the estimates at each iteration. The subset is chosen stochastically, so that each subset has a probability of being chosen equal to $1/B^r$. This algorithm can be seen as an equivalent of the Stochastic Average Gradient algorithm[7] that uses Jensen surrogates instead of Lipschitz quadratic surrogates for the X-Ray CT application. Also, it is important to note that we only use the previous estimate information for the data-fitting term; a similar approach can be used for the regularization term as well, but this would require the storage of B^r images. Not only does this make the one-dimensional optimization to update the estimate slower, but also increases the storage requirement. Experimentally, we have not found any advantages and will not include it here.

Algorithm 3 presents the Stochastic Average Jensen Surrogates Optimization for X-Ray CT (SA-JS). For a very large number of subsets, the requirement of storing B^r back projected estimates in the image domain can be problematic if the memory capacity of the computing architecture being used is low. As seen in the image update, at each iteration we use the most recent sum of the back projected estimates. In a naive implementation, this computation would add $\mathcal{O}(B^r N)$ to the computational complexity. A better approach is to store the most recent sum as another variable, and when a subset is chosen, subtract the old back projected estimate of the corresponding estimate, and add the newly computed one to the sum. This version would require $\mathcal{O}(2N)$ in additional computations and an additional $\mathcal{O}(N)$ of storage, which is more reasonable than the naive implementation.

3.3 Ordered Subsets Average Jensen Surrogates Optimization for X-Ray CT

In addition to the stochastic choice of subsets presented in the previous section, one can also choose them in a cyclic manner as in the ordered subsets method. The resultant algorithm, called Ordered Subsets Average Jensen Surrogates Optimization for X-Ray CT (OSA-JS), is shown in Algorithm 4.

Algorithm 3 Stochastic Average Jensen Surrogates Optimization for X-Ray CT (SA-JS)

Input : $\boldsymbol{x}^{(0)} \in \mathbb{R}_+^N, \boldsymbol{d}, \boldsymbol{I_0} \in \mathbb{R}_+^M, \boldsymbol{H} \in \mathbb{R}^{M\times N}, \lambda \geq 0, \delta > 0$, $\mathcal{B}_k^r$ for $k = 0, 1, ..., (B^r - 1)$, $\boldsymbol{x}^{(0,k)} \in \mathbb{R}_+^N$ for $k = 0, 1, ..., (B^r - 1)$

Pre-compute $b_j^k = \sum_{i\in\mathcal{B}_k^r} d_i h_{ij}$, $\forall j, k$.

Pre-compute $Z = \max_i \sum_j h_{ij}$.

for $n = 0, 1, 2, ...$ **do**

  Choose k from $\{0, 1, ..., (B^r - 1)\}$ randomly.

  $\boldsymbol{x}^{(n,k)} = \boldsymbol{x}^{(n)}$

  $q_i^{(n)} = I_{0,i} \exp(-\sum_j h_{ij} x_j^{(n)})$, $\forall i \in \mathcal{B}_k^r$.

  $b_j^{(n,k)} = \sum_{i\in\mathcal{B}_k^r} q_i^{(n)} h_{ij}$, $\forall j$.

  $x_j^{(n+1)} = \text{argmin}_{x\geq 0} \sum_k b_j^k (x - x_j^{(n)}) + \sum_k b_j^{(n,k)} / Z \exp(-Z(x - x_j^{(n)})) + \lambda B_{\boldsymbol{s}}(\boldsymbol{x}; \boldsymbol{x}^{(n)})$, $\forall j$.

end

Algorithm 4 Ordered Subsets Average Jensen Surrogates Optimization for X-Ray CT (OSA-JS)

Input : $\boldsymbol{x}^{(0)} \in \mathbb{R}_+^N, \boldsymbol{d}, \boldsymbol{I_0} \in \mathbb{R}_+^M, \boldsymbol{H} \in \mathbb{R}^{M\times N}, \lambda \geq 0, \delta > 0$, $\mathcal{B}_k^r$ for $k = 0, 1, ..., (B^r - 1)$, $\boldsymbol{x}^{(0,k)} \in \mathbb{R}_+^N$ for $k = 0, 1, ..., (B^r - 1)$

Pre-compute $b_j^k = \sum_{i\in\mathcal{B}_k^r} d_i h_{ij}$, $\forall j, k$.

Pre-compute $Z = \max_i \sum_j h_{ij}$.

for $n = 0, 1, 2, ...$ **do**

  $k =$ mod (n, B^r)

  $\boldsymbol{x}^{(n,k)} = \boldsymbol{x}^{(n)}$

  $q_i^{(n)} = I_{0,i} \exp(-\sum_j h_{ij} x_j^{(n)})$, $\forall i \in \mathcal{B}_k^r$.

  $b_j^{(n,k)} = \sum_{i\in\mathcal{B}_k^r} q_i^{(n)} h_{ij}$, $\forall j$.

  $x_j^{(n+1)} = \text{argmin}_{x\geq 0} \sum_k b_j^k (x - x_j^{(n)}) + \sum_k b_j^{(n,k)} / Z \exp(-Z(x - x_j^{(n)})) + \lambda B_{\boldsymbol{s}}(\boldsymbol{x}; \boldsymbol{x}^{(n)})$, $\forall j$.

end

4. RESULTS

In this section, we compare the proposed algorithms and the gradient descent variants using real data. Before presenting the results, we briefly explain the competing algorithms implemented.

The gradient descent method is a popular technique for iterative convex minimization. In an iteration, one forms a surrogate function around the current estimate that is equal to

$$\Phi(\boldsymbol{x}^{(n)}) + \nabla(\Phi(\boldsymbol{x}^{(n)}))^T(\boldsymbol{x} - \boldsymbol{x}^{(n)}) + \frac{\alpha^{(n)}}{2}\|\boldsymbol{x} - \boldsymbol{x}^{(n)}\|_2^2. \tag{9}$$

When the function of interest is twice differentiable and Lipschitz continuous, this method has a guarantee of convergence when the quadratic scaling factor $\alpha^{(n)}$ is chosen to be equal to L_Φ, where L_Φ is the Lipschitz gradient constant of Φ.[‡] For our case, the gradient is equal to

$$\begin{aligned} \nabla\Phi(\hat{\boldsymbol{x}}) &= \nabla f(\hat{\boldsymbol{x}}) + \lambda\nabla\beta(\hat{\boldsymbol{x}}) & (10)\\ &= \boldsymbol{H}^T(\boldsymbol{d} - \hat{\boldsymbol{q}}) + \lambda\boldsymbol{C}^T(\nabla\tilde{\beta}(\boldsymbol{C}\hat{\boldsymbol{x}})), & (11)\end{aligned}$$

where $\hat{q}_i = I_{0,i} \exp(-\sum_j h_{ij}\hat{x}_j)$. The Lipschitz gradient constant L_Φ is computed using the following inequality

[‡]Using an arbitrary initial $\alpha^{(n)}$ and decreasing it when the objective function starts increasing is a popular technique when this constant is not computable. This method, backtracking, can also be used with Jensen surrogates. For simplicity, this investigation is left as future work.

on the Hessian,

$$\begin{aligned}\nabla^2\Phi(\boldsymbol{x}) &= \nabla^2 f(\boldsymbol{x}) + \lambda\nabla^2\beta(\boldsymbol{x}) && (12)\\ &\leq \max_{l,i}\frac{\partial^2\tilde{f}_i(l)}{\partial l^2}\boldsymbol{H}^T\boldsymbol{H} + \lambda\max_{l,k}\frac{\partial^2\tilde{\beta}_k(l)}{\partial l^2}\boldsymbol{C}^T\boldsymbol{C} && (13)\\ &\leq \max_i I_{0,i}\boldsymbol{H}^T\boldsymbol{H} + \lambda\boldsymbol{C}^T\boldsymbol{C}, && (14)\end{aligned}$$

which in return can be used to find L_Φ as

$$L_\Phi = \Lambda_{max}(\max_i I_{0,i}\boldsymbol{H}^T\boldsymbol{H} + \lambda\boldsymbol{C}^T\boldsymbol{C}) \geq \max_{\boldsymbol{x}}\Lambda_{max}(\nabla^2\Phi(\boldsymbol{x})), \quad (15)$$

where $\Lambda_{max}(\cdot)$ is the operator that returns the maximum eigenvalue of the given square matrix. When the system matrix is scaled down in the image domain, we experimentally found that $\Lambda_{max}(\boldsymbol{H}^T\boldsymbol{H})$ follows a linear relationship with the fully sampled version. Thus, we computed it by finding $\Lambda_{max}(\boldsymbol{H}^T\boldsymbol{H})$ for a scaled-down system matrix and then multiplying by the appropriate scale factor. For the regularization matrix part, we computed $\Lambda_{max}(\boldsymbol{C}^T\boldsymbol{C})$ for a subset of the image volume, 5 z-slices, to make it computationally tractable.

Algorithm 5 presents the Gradient Descent Optimization for X-Ray CT (Full-GD). For convenience, we will compare the ordered subset (OS-GD) and the stochastic average (SA-GD) variants of the gradient descent method, which are straightforward extensions of the base case algorithm.

Algorithm 5 Gradient Descent Optimization for X-Ray CT (Full-GD)

Input : $\boldsymbol{x}^{(0)} \in \mathbb{R}_+^N, \boldsymbol{d}, \boldsymbol{I_0} \in \mathbb{R}_+^M, \boldsymbol{H} \in \mathbb{R}^{M\times N}, \lambda \geq 0, \delta > 0$.
Pre-compute $b_j = \sum_i d_i h_{ij}$, $\forall j$.
for $n = 0, 1, 2, ...$ **do**
  $q_i^{(n)} = I_{0,i}\exp(-\sum_j h_{ij}x_j^{(n)})$, $\forall i$.
  $b_j^{(n)} = \sum_i q_i^{(n)} h_{ij}$, $\forall j$.
  $x_j^{(n+1)} = \left[x_j^{(n)} - \frac{1}{L_\phi}\left((b_j - b_j^{(n)}) + \lambda\frac{\partial\beta(\boldsymbol{x}^{(n)})}{\partial x_j^{(n)}}\right)\right]_+$, $\forall j$.
end

The acceleration methods have been investigated using two real-data scans acquired on a SureScan™ $x1000$ Explosive Detection System. For all the methods presented, the regularization parameters $\lambda = 15000$ and $\delta = 0.001$ were used. All reconstructed image volumes consist of 240 rows, 400 columns, and a varying number of slices along the z-axis. In order to compare the rates of convergence, we look at the normalized function errors $(\Phi(\boldsymbol{x}^{(n)}) - \Phi(\boldsymbol{x}^*))/\|\Phi(\boldsymbol{x}^*)\|$, where the objective function value at the optimum was computationally found by running the algorithms for many more iterations. We examine the performance of our algorithms with a relatively low number of subsets, 8, and a large number of subsets, 64.

We compare the proposed algorithms SA-JS and OSA-JS with

- the full variant of Jensen surrogates, Full-JS,
- the momentum and varying step-size variant, Fast-JS,[6]
- the ordered subsets with Jensen Surrogates method, OS-JS,
- the Stochastic Average Gradient method,[7] SA-GD, and
- the ordered subsets with gradient descent method, OS-GD.

Figures 1-3 show three slices each of axial, sagittal, and coronal views of the reconstructed volume for Bag #1. Figure 4 presents the normalized function errors for different algorithms vs. the number of effective data

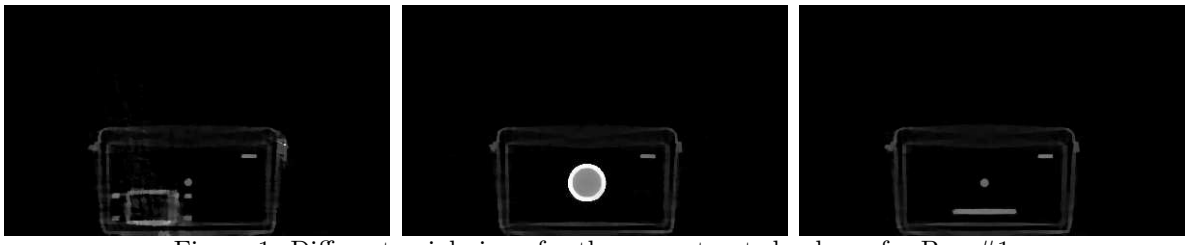

Figure 1: Different axial views for the reconstructed volume for Bag #1.

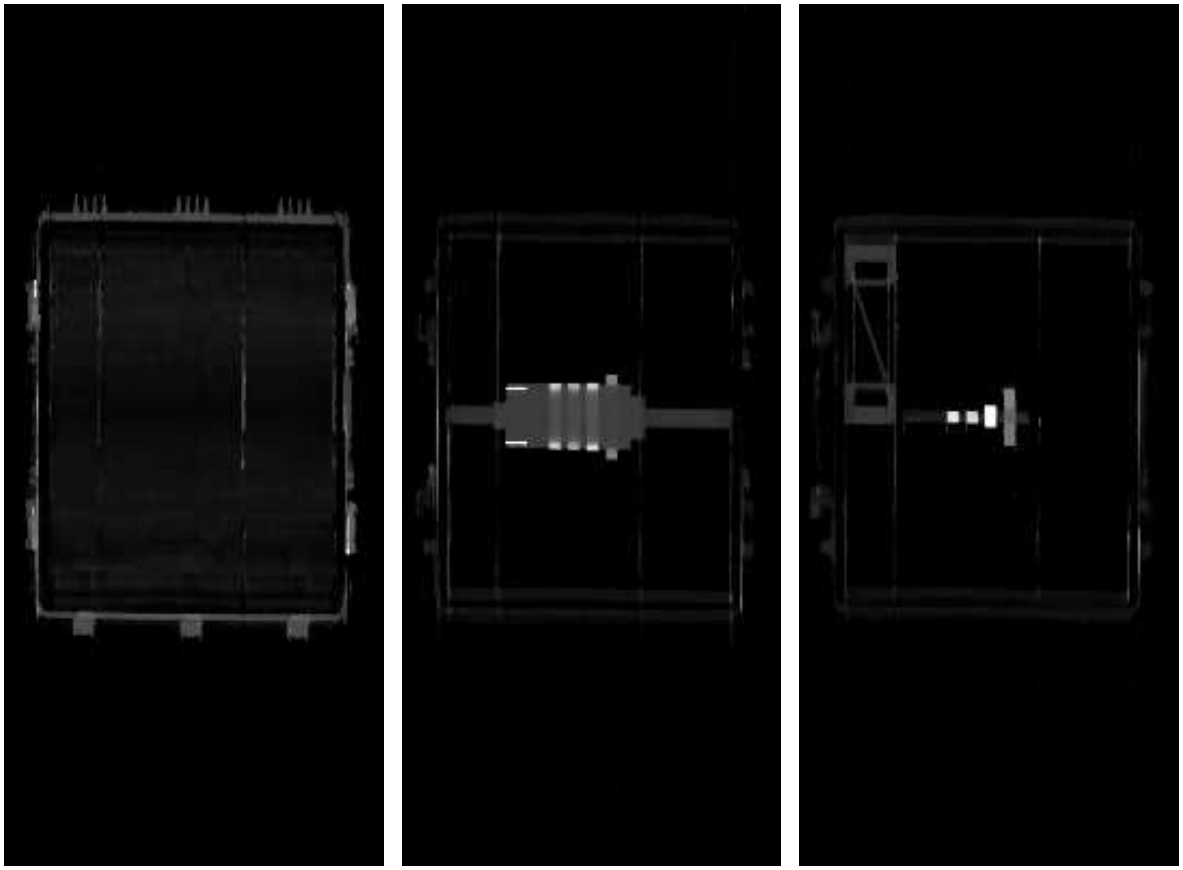

Figure 2: Different sagittal views for the reconstructed volume for Bag #1.

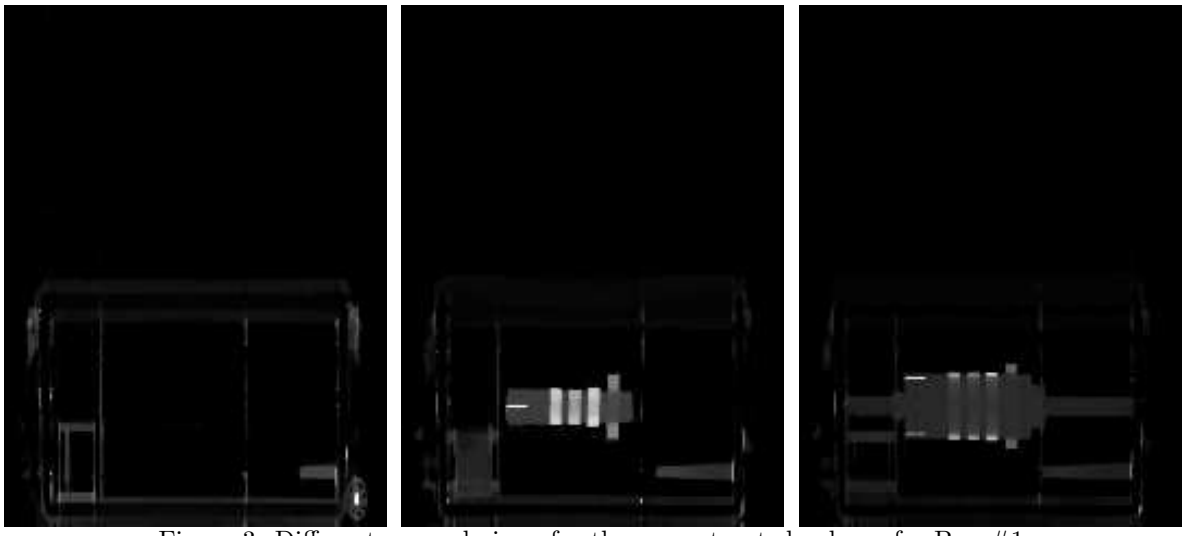

Figure 3: Different coronal views for the reconstructed volume for Bag #1.

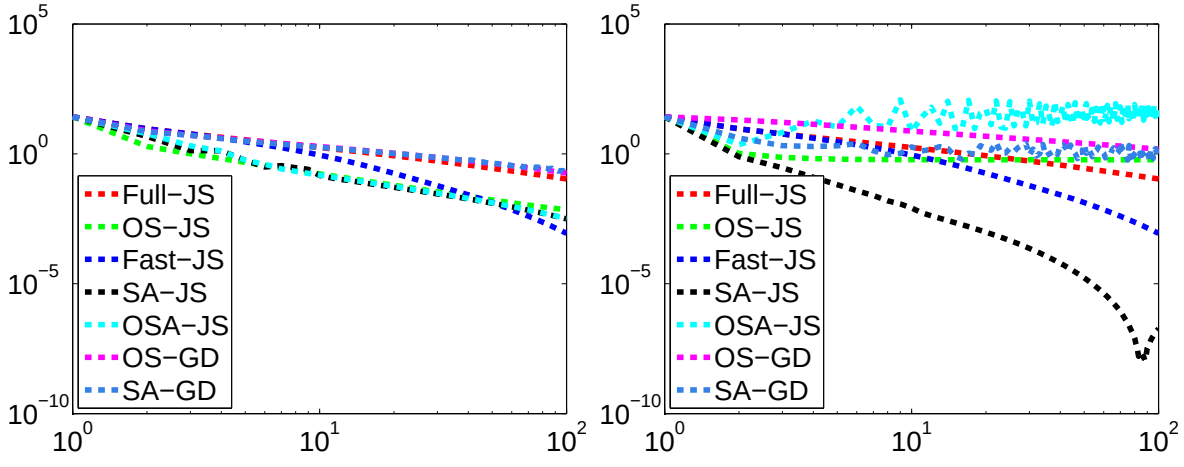

(a) 8 Subsets (b) 64 Subsets

Figure 4: Normalized function errors vs. number of effective data passes for Bag #1. (a) shows the results for 8 subsets, while (b) shows the results for 64 subsets.

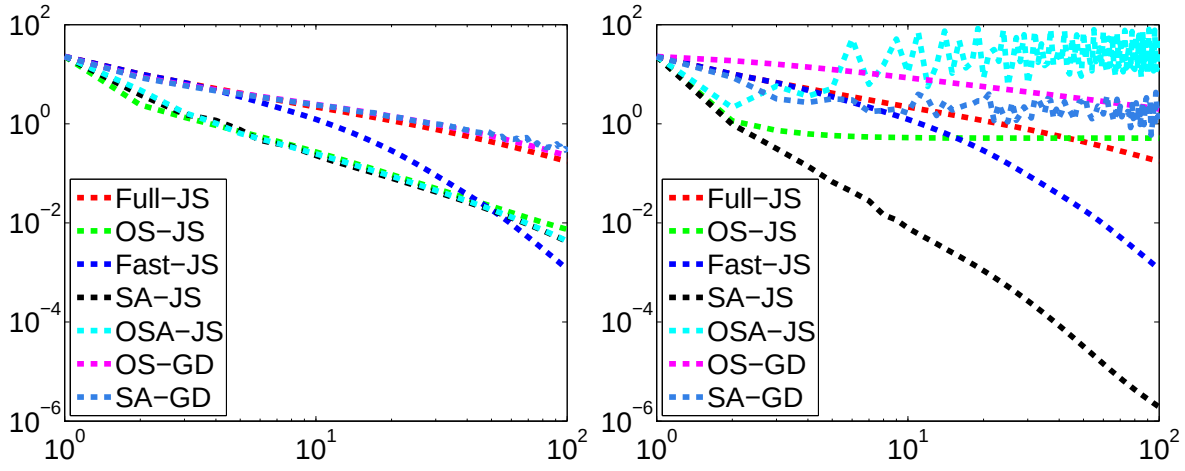

(a) 8 Subsets (b) 64 Subsets

Figure 5: Normalized function errors vs. number of effective data passes for Bag #2. (a) shows the results for 8 subsets, while (b) shows the results for 64 subsets.

passes for Bag #1 for two subset settings while Figure 5 presents the errors for Bag #2. As seen in the figures, for 8 subsets, the Fast-JS algorithm slightly outperforms the SA-JS and OSA-JS algorithms proposed. For 64 subsets, we see that SA-JS algorithm we proposed outperforms the other competing algorithms.

It is also important to note that the proposed Jensen surrogate algorithms outperform their gradient descent equivalents for all cases. Another observation is that the cyclic choice we proposed for the average scheme we proposed works competitively for the 8 subsets case, but performs worse than its stochastic counterpart for 64 subsets.

5. CONCLUSION

In this paper, we investigated incremental methods and proposed one stochastic and one deterministic algorithm that use Jensen surrogates and compared these methods with other competing algorithms using data collected from two bags by a baggage scanner. The Stochastic Average variant (SA-JS) we proposed performs competitively for the small number of subsets and outperforms the other competing algorithms for the large number of subsets in terms of number of effective data passes. This is an indicator that it should perform better as the number of subsets increases. However, implementation of the larger number of subsets requires additional storage and more image updates per effective data pass, which both add computational costs. Thus, we plan to conduct a thorough time analysis of the proposed algorithms as well as to explore the "hybrid" variants.

REFERENCES

1. D. P. Bertsekas, "Incremental gradient, subgradient, and proximal methods for convex optimization: A survey," *Optimization for Machine Learning* **2010**, pp. 1–38, 2011.
2. J. A. O'Sullivan and J. Benac, "Alternating minimization algorithms for transmission tomography," *Medical Imaging, IEEE Transactions on* **26**(3), pp. 283–297, 2007.
3. S. Degirmenci, D. G. Politte, C. Bosch, N. Tricha, and J. A. O'Sullivan, "Acceleration of iterative image reconstruction for x-ray imaging for security applications," in *IS&T/SPIE Electronic Imaging*, pp. 94010C–94010C, International Society for Optics and Photonics, 2015.
4. K. Lange, J. Fessler, *et al.*, "Globally convergent algorithms for maximum a posteriori transmission tomography," *Image Processing, IEEE Transactions on* **4**(10), pp. 1430–1438, 1995.
5. J. Nocedal and S. Wright, *Numerical optimization*, Springer Science & Business Media, 2006.
6. S. Degirmenci, *A General Framework of Large-Scale Convex Optimization Using Jensen Surrogates and Acceleration Techniques*. PhD thesis, Washington University in St. Louis, 2016.

7. N. L. Roux, M. Schmidt, and F. R. Bach, "A stochastic gradient method with an exponential convergence rate for finite training sets," in *Advances in Neural Information Processing Systems*, pp. 2663–2671, 2012.